QORINAI

CORPORATE PROFILE · 2026

NVIDIA HGX B300 configured and tested in Australia · GB300 NVL72 halls in design

Australian AI infrastructure, built from the grid to the GPU.

QORINAI develops, builds and operates AI data centres — land and power, design and construction, hardware, cluster deployment, 24×7 GPU operations and a sovereign cloud on top. Every link of the chain, one accountable team.

26 MW+
delivered by our delivery partners

150 kW
per rack, liquid-cooled design

GW-class
campus in masterplanning

100 %
Australian owned & operated

ABOUT QORINAI

QOR is the core. IN is the intelligence layer. AI is what runs on top.

QORINAI is an Australian-owned developer and operator of AI data centres. We exist because the gigawatt era of AI infrastructure is a power project, a construction project, a hardware project and a software operation at once — and Australia had no single company organised to deliver all of it.

QOR

The core

The chip etched in our symbol, the hall, the substation: the physical heart of the business, and the part we build with our own hands.

IN

The intelligence layer

The fabric, the telemetry, the operations runbooks, the people on site at 2 a.m. — everything that connects the core.

AI

What runs on top

The models, the clusters, the sovereign cloud our customers rent by the GPU-hour.

MISSION

Energisable land, into sovereign compute.

Australia's data-centre pipeline is measured in gigawatts; the capacity that will actually be energised this decade is a fraction of it. Our job is to close that gap — securing power, building halls at GPU density, and running the fleets inside them under Australian control.

HOW WE WORK

- ✓ **Power first.** The grid process runs before the design process — a hall without a connection is a shed.
- ✓ **Owner-managed delivery.** Specialist trades engaged directly; no master EPC. The margin and know-how stay with the project.
- ✓ **Density by design.** Every hall specified for 150 kW liquid-cooled racks from day one.
- ✓ **Sovereign by default.** Australian-owned, -staffed and -hosted. No training on customer data.

WHY QORINAI

Australia has no shortage of planned data centres. It has a shortage of energisable land, liquid-cooled halls and people who can run a GPU fleet.

So we built a company around all three.

POWER FIRST

The connection before the concrete

Transmission and distribution enquiries, system-strength studies and three supply layers — grid, interim distribution, on-site generation — before a hall is designed.

DENSITY BY DESIGN

150 kW liquid-cooled racks from day one

Direct-to-chip cooling, rear-door heat exchange and air halls in every campus, so the next NVIDIA or AMD generation is a rack swap, not a rebuild.

OWNER-MANAGED

No master EPC in the middle

Electrical, mechanical, liquid-cooling and compliance trades engaged directly. The margin and the know-how stay with the project.

SOVEREIGN

Australian owned, staffed and hosted

No prompt caching, no training on your data, no offshore approval between you and your hardware.

SUPPLY CERTAINTY

A manufacturer on the share register

Hyperscalers, Australia's open-OEM manufacturer, is our strategic shareholder — configured and tested in Australia, Australian warranty, lab test drives.

ONE OPERATOR

The team that designed the loop answers the 2 a.m. alert

One contract, one SLA, one phone number from build through operations.

THE AIDC CHAIN

Eight links from the transmission line to the token.

Most AI projects stall between the substation and the GPU — four companies, four contracts, nobody accountable for the gap. QORINAI is organised so that the gap doesn't exist.

01

Land & power

Sites beside EHV corridors, formal connection enquiries, gas firming and battery storage.

4 MVA → 1 GW

02

Design

Tier III+ MEP, three cooling tiers, 220 / 33 kV reticulation, design PUE ≤ 1.2.

150 KW PER RACK

03

Build

Transportable 2.5 MW modules, owner-managed subcontracting, phased power blocks.

10–50 MW BLOCKS

04

Hardware

HGX B300, GB300 NVL72 and AMD Instinct systems configured and tested in Australia via Hyperscalers.

AUSTRALIAN-BUILT · AU WARRANTY

05

Deploy

Rack, cable, InfiniBand XDR fabrics, NCCL / HPL burn-in and thermal soak.

HANDED OVER BENCHMARKED

06

Operate

24×7 GPU fleet management, firmware lifecycle, RMA with hot spares, monthly reporting.

99.9 % NODE SLA

07

Host

Capacity hosting and compute leasing in operating Tier III+ halls — no greenfield build to wait for.

30–150 KW PER RACK · LIVE TODAY

08

Compute

Sovereign GPU cloud, 100+ open-weight models, private cloud — billed in AUD.

NO TRAINING ON YOUR DATA

LINKS 01–02 · LAND & POWER, DESIGN

The power process runs first. Halls are designed for the rack you'll deploy next.

Only a fraction of Australia's planned capacity will be energised this decade. The binding constraint is a formal, rules-based connection — so that is the first thing we secure.

THREE SUPPLY LAYERS THAT BACK EACH OTHER UP

Transmission

NER connection enquiries; new terminal-station or 220 kV options; PSS/E · PSCAD studies.

Interim distribution

Distribution pathway for construction and early IT load while transmission runs.

On-site generation

Gas firming, grid-forming BESS and behind-the-meter PPAs so the campus ramps on schedule.

THREE COOLING TIERS IN EVERY CAMPUS · TIER III+ · DESIGN PUE ≤ 1.2

Tier A · Air

≤ 30 kW / rack

Hot-aisle containment for storage, networking, PCIe GPU and CPU nodes.

Tier B · RDHx

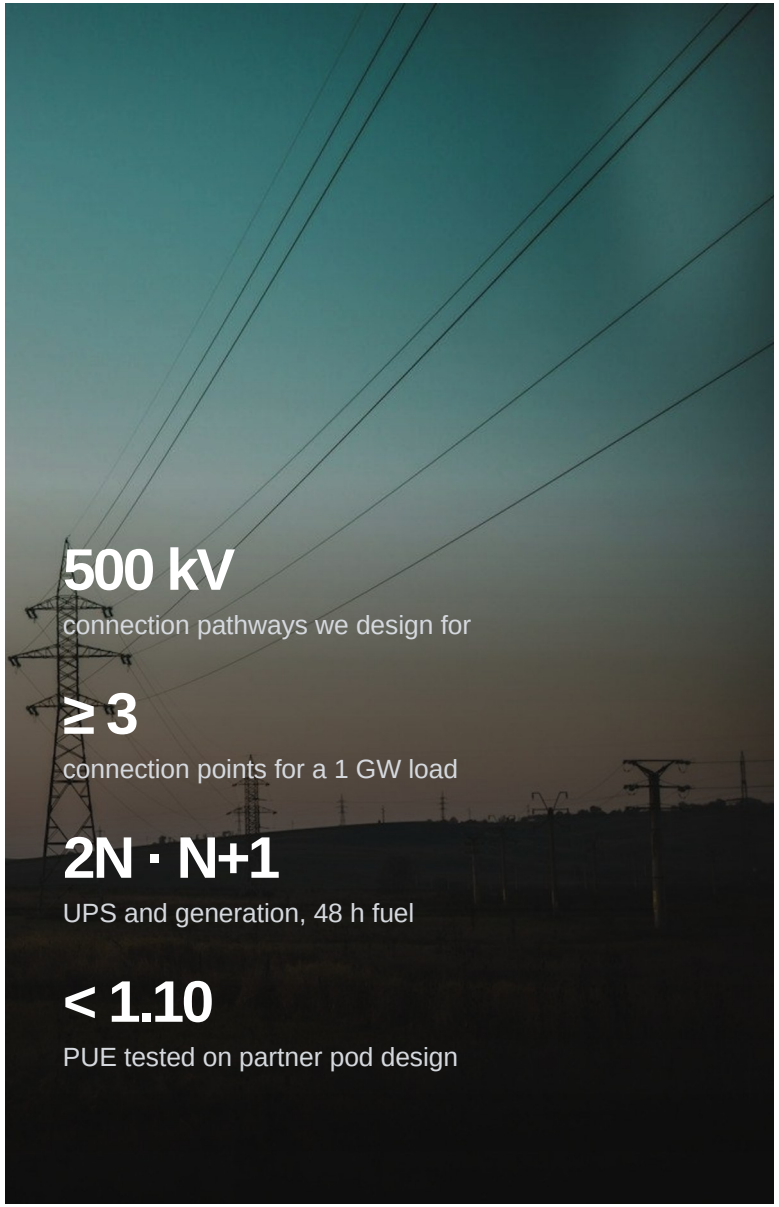
30–80 kW / rack

Rear-door heat exchange for HGX B200 / B300 air-cooled nodes without changing the server.

Tier C · DLC

80–150 kW / rack

Direct-to-chip liquid, CDU per row, 415 V busway. Built for GB300 NVL72 and beyond.



Modular blocks, owner-managed delivery. Handed over benchmarked, not boxed.

Transportable 2.5 MW modules

Factory-built, self-contained units with their own distribution, cooling and security envelope. Eight modules deliver 20 MW; each can be dismantled, refurbished or redeployed.

SIGNED → FIRST POWER IN DAYS

Modular launch on existing supply

Where a site has distribution capacity today, a 2.5–3 MW modular launch goes live on a 14-month programme while the transmission connection is in process.

14-MONTH PROGRAMME

Campus power blocks

Repeatable 10–50 MW blocks, each with its own substation capacity, cooling plant and data hall — capital committed in phases, tenants live while later blocks are in the ground.

10–50 MW BLOCKS

CONVENTIONAL · MASTER EPC

- 10–20 % EPC margin on every trade
- Interface costs between owner, EPC and subcontractors
- Liquid-cooling iterations wait for variations
- Know-how stays with the EPC

QORINAI MODEL · OWNER-MANAGED

- Electrical, mechanical / liquid cooling, cabling and compliance trades engaged directly
- No master-EPC margin or interface cost
- Design changes resolved on site, same week
- Know-how accumulates in our runbooks and product designs
- Proven subcontractor bench re-used on the next cluster

LINK 05 · DEPLOY

5.1 Rack & power

Rack elevations, busway tap-offs, A+B PDU balancing, liquid manifolds for DLC nodes.

5.2 Fabric build

Rail-optimised InfiniBand XDR or 800G RoCEv2, 1:1 non-blocking, dual-routed inter-hall links.

5.3 Burn-in & validation

Firmware baseline, DCGM, NCCL all-reduce, HPL, thermal soak — documented per node.

5.4 Platform handover

Slurm or Kubernetes, storage mounts, tenant isolation, monitoring and runbooks on day one.

GPUs are only an asset when they are busy.

A B300 node that sits idle, throttles on a warm coolant loop or waits three weeks for an RMA is a depreciating liability. We run GPU fleets the way a utility runs a grid: measured continuously, maintained before failure, reported monthly to the owner.

99.9 %
node availability, monthly SLA

< 15 min
P1 response, on-site engineer

< 4 h
general smart hands

30–150 kW
per rack, dual A+B feeds

06 · OPERATE — MEASURED ON EVERY GPU

- SM occupancy and tensor-core utilisation
- HBM temperature, ECC and row-remap counts
- NVLink / NVSwitch link errors and retrains
- Xid faults, driver resets, power per GPU and rack
- Coolant supply temperature and ΔT per CDU
- Fabric health, job throughput and queue depth

24×7 NOC · firmware, driver and CUDA lifecycle · RMA with hot spares · Slurm / Kubernetes · utilisation reporting for owners and lenders

07 · HOST — COLOCATION AT GPU DENSITY

Footprints Quarter rack to private cages and suites

Power & cooling Dual-feed standard, up to 150 kW per rack on DLC halls, redundant plant

Connectivity 40 Gbps+ cross-connects, dark fibre, IP transit, major AU and international carriers

Smart hands < 15 min P1, < 4 h general — Australian engineers on site

08 · COMPUTE — SOVEREIGN AI & CLOUD

- GPU-as-a-Service — reserved pools, AUD billing
- GPU & CPU clusters; single GPU to full nodes
- 100+ open-weight models, OpenAI-compatible API
- Bare metal, VMs, private OpenStack cloud
- NVMe block and S3 object storage, DR & backup
- HPC and burst VFX rendering

Hosted in Australia · billed in AUD · no prompt caching · no training on your data

GPU HOSTING & COMPUTE LEASING · LIVE TODAY

GPUs go live in halls that already exist.

Campuses take years; the practical near-term GPU supply in Australia is what can be hosted in existing, powered, Tier III+ halls. QORINAI deploys incremental GPU clusters into built, operating partner facilities and runs them under term contracts — no greenfield risk, no waiting for a grid connection.

Tier III+
operating, powered halls — no greenfield build

128 nodes
the cluster unit we deploy — 1,024 GPUs, 800G fabric

30–150 kW
per rack, air, RDHx or direct-to-chip liquid

99.9 %
node availability, 24×7 Australian NOC

YOU OWN THE GPUS

We host and operate them

Capacity hosting for GPU owners, neoclouds and enterprise AI teams: power, cooling, network, security and smart hands in one monthly AUD fee. Fleet management adds 24×7 monitoring, firmware lifecycle, RMA with hot spares and reporting for owners and lenders.

MONTHLY AUD · PER RACK OR PER GPU

YOU NEED COMPUTE

Lease it from us on a term

Reserved HGX B300 nodes or full 128-node clusters at a fixed per-GPU-hour or monthly price for the term — one to five years, take-or-pay, no escalation. Hosted in Australia, no training on your data. Start with part of a cluster; scale to a dedicated hall.

FIXED AUD PRICE · 1–5 YEAR TERMS

YOU WANT THE YIELD

Co-fund a contracted cluster

GPU clusters are a second asset layer on top of the real estate. Investors and lenders fund hardware configured and tested in Australia by Hyperscalers, deployed into an operating facility and fully contracted to a lessee before the pallets arrive; QORINAI deploys, operates and reports.

CONTRACTED BEFORE DEPLOYMENT

HOW A CLUSTER GOES LIVE

1. Hardware allocated and configured through Hyperscalers
2. Staged into an operating Tier III+ hall — liquid-cooled racks ready
3. Burn-in: NCCL, HPL, thermal soak — handed over benchmarked
4. 24×7 operations, 99.9 % SLA, monthly reporting
5. Move to a dedicated hall or the campus later — no re-platforming

Hosting facilities are operating partner sites; counterparties and terms under NDA

One project, one vehicle. Every partnership ring-fenced.

Each partnership — a GPU cluster and its lease, a data-centre build, a new country — is its own special-purpose vehicle under the Qinari name, with its own assets, contracts, capital, lenders and returns. QORINAI operates every vehicle; partners invest only in the one they choose.

Qinari Compute

GPU equipment · ownership & lease

Owns the cluster. Leases the compute.

- Holds: servers, fabric, spares, warranties
- Contracts: capacity or equipment lease; O&M with QORINAI
- Funded by: equipment finance and partner equity

EQUIPMENT FINANCIERS · GPU OWNERS · OFFTAKE CUSTOMERS

Qinari Infrastructure

Data-centre development & construction

Holds the land, the power and the hall.

- Holds: land, connection rights, plant, buildings
- Contracts: development, construction, hosting or lease; O&M
- Funded by: partner equity, project or construction debt

ENERGY COMPANIES · INFRASTRUCTURE FUNDS · EQUIPMENT-FOR-EQUITY SUPPLIERS

Qinari International

Vehicles outside Australia

The same model, incorporated where the partner is.

- Holds: compute, infrastructure or both — as the local project needs
- Contracts: local partner agreement, QORINAI services, hardware supply
- Funded by: local partner capital, export or development finance

JAPAN · SINGAPORE · MALAYSIA · MIDDLE EAST

WHY A VEHICLE PER PROJECT

Ring-fencing a delay or lease event on one asset cannot touch another vehicle

Clean cap table partners invest at project level, in the vehicle they choose

Lender-ready single-purpose, asset-backed; no consolidation, no cross-default

Alignment QORINAI co-invests and operates each SPV under a services agreement

Jurisdiction fit an Australian Pty Ltd for Australian assets; a local vehicle abroad

Exit optionality refinance, hold or sell — one vehicle at a time

SITS IN THE QINARI SPV

Legal title to the equipment, land or building · customer lease, capacity or hosting contract · project debt and security over the vehicle's own assets · partner shareholdings and the shareholders' agreement · ring-fenced accounts and monthly reporting

STAYS WITH QORINAI

Engineering, deployment and 24×7 operations teams · the IntraLLM platform, methods and designs · the hardware supply relationship through Hyperscalers · development-management, O&M and services agreements with each vehicle · a co-investment in each vehicle alongside partners

Qinari — "core in fabric"; general approach, not an offer; terms agreed per vehicle under NDA

The right accelerator for every AI workload.

Configured, tested and warranted in Australia through our strategic shareholder Hyperscalers — manufactured in Australia, full Australian warranty, channel pricing on request.



FEATURED · 8U

NVIDIA HGX B300 8-GPU

Eight Blackwell Ultra GPUs on one NVLink domain, dual Intel Xeon 6, 800G fabric — delivered fully assembled and tested in Australia.

GS7 B300 · G894-SD3-AAX7



RACK-SCALE

NVIDIA GB300 NVL72 · B300 density

72 GPUs and 36 Grace CPUs on one NVLink domain, ~21 TB HBM3e, ~130 kW liquid-cooled — or 72 B300 per rack on 3U density nodes.

DLC · 415 V BUSWAY



INFERENCE · 7U

AMD Instinct MI325X 8-GPU

8× MI325X, 256 GB HBM3e each — 2 TB per node on EPYC 9005 Turin. Memory-bound inference of large open-weight models.

S8MT · D75T-7U



FINE-TUNING · 8U

NVIDIA HGX B200 8-GPU

Blackwell at an established price point; 8× 400 Gb/s NDR, ~14 kW node, air-cooled or DLC-ready.

GX5 B200 · G893-SD1-AAX5



FABRIC

InfiniBand XDR & 800G Ethernet

Quantum-X800, Spectrum-X, BlueField-3 DPUs, ONIE bare-metal switches 1–400G, 800G optics — configured and tested in Australia.

1G → 800G



STORAGE

NVMe all-flash, JBOD, parallel FS

Gen5 NVMe all-flash nodes, 4U JBODs 42–102 bays, Lustre / WEKA-class parallel file systems that keep 8-GPU nodes fed.

NVME · JBOD · PFS

LINK 04 · HARDWARE · FEATURED NODE

NVIDIA HGX B300 8-GPU (G894-SD3) — our standard training and inference node.



G894-SD3 · XEON 6 · 8U DP · HGX B300

Eight Blackwell Ultra GPUs on one NVLink domain, dual Intel Xeon 6, 800G fabric — delivered fully assembled and tested in Australia.

GPU	8× NVIDIA B300 Blackwell Ultra · NVLink + NVSwitch
CPU	2× Intel Xeon 6 6767P · 64 cores each
FABRIC	ConnectX-8 · 8× 800 Gb/s OSFP
FORM	8U · air-cooled or direct-to-chip liquid
DELIVERY	Level 10 assembled & tested · Australian warranty

Full Australian warranty

Configured and burned in before delivery

Test-drive in HyperLabs before you buy

Configured & tested in Australia · Canberra

EVERYTHING ELSE IN THE RACK

The complete open-OEM range through our shareholder — configured, tested and warranted in Australia.

Servers GPU (HGX, PCIe, Instinct), PCIe Gen5 Xeon 6 / EPYC 9005, 1U–10U

Hyper-converged 4-node 2U, multi-node SSD, SDS / SDC, micro servers

Storage 4U JBODs, NVMe all-flash, SDS-ready storage servers

Switches T1000 / T3000 / T5000 / T7000 bare-metal ONIE, 1G–400G

OCP hyperscale Rackgo M / X / X-RSD open-rack compute and storage

Edge & devices PCIe Gen5 workstations, far-edge, telco / 5G edge

Commodities CPU · RAM · SSD · NVMe · NICs · optics · GPUs

Lab Lab-as-a-Service test drives, IP-appliance design

STRATEGIC SHAREHOLDER

Backed by Australia's open-OEM server manufacturer.

HYPERSCALERS

Hyperscalers has designed, manufactured and supported open-OEM servers, storage and networking in Australia since 2014 — the hardware behind cloud providers, telcos and enterprise IT. As our strategic shareholder it brings local manufacturing, configuration and testing in Australia, a full Australian warranty and a lab to test-drive configurations before purchase.



EST. 2014 · CANBERRA | MADE IN AUSTRALIA | FULL AUS WARRANTY | HYPERLABS TEST-DRIVE | AU & US OPERATIONS

WHAT A MANUFACTURER ON THE SHARE REGISTER CHANGES

Supply certainty

Allocation of NVIDIA HGX B300 and AMD Instinct systems through Hyperscalers, configured and burned in in Australia before the pallet reaches the hall.

Sovereign warranty

Australian warranty, spares and engineers. A failed GPU tray is swapped from hot spares and RMA'd domestically.

Open economics

Open-OEM pricing and generic commodities keep capital cost per GPU — and the hosting price per GPU-hour — below Tier 1 OEM builds.

Shared customer base

Cloud providers, colocation operators, enterprise IT, telcos and consultants — a partner who buys from Hyperscalers can host with QORINAI.

Test before you commit

Every QORINAI reference configuration can be test-driven in HyperLabs — the cluster you benchmark is the cluster we deploy.

SOFTWARE · INTRALLM · QORINAI PROPRIETARY PLATFORM

Your own AI, inside your own network.

IntraLLM is QORINAI's turn-key enterprise-AI platform: run and train your own LLM and multimodal models on-premises, in a private tenancy in our halls, or on the sovereign cloud — hardware, orchestration, models, assistants and automation in one console. Your data never leaves the building.



ON-PREM OR PRIVATE CLOUD

AIR-GAPPED CAPABLE

BRING YOUR OWN MODEL

NVIDIA & AMD GPUS

Business Assistant

Internal and customer-facing assistants grounded in your documents, systems and policies.

Process Automation

Agents that read screens, fill forms and complete workflows across ERP, CRM and ticketing.

Business Assurance

Quality-assure construction documentation, cyber posture and marketing output against your standards.

Bring Your Own Model

Load, fine-tune and serve proprietary or fine-tuned models alongside the curated repository.

Model Repository

Curated open-weight LLM / LMM models — GLM-4.6V, Qwen 2.5, DeepSeek-class reasoning — versioned locally.

Cluster Manager

GPU job orchestration and live telemetry — utilisation, memory, power, heat — across every node.

SQL Agent

Plain-English questions become optimised, explained SQL on SQLite, MySQL, PostgreSQL, Oracle and SQL Server.

Support

Australian engineering support for the platform, the models and the hardware underneath — one number.

Developed with our strategic shareholder Hyperscalers · licensed per deployment

Every enterprise will need to train and mature its own AI models. The question is whether that knowledge stays inside the company.

IP retention

Models trained inside the network turn staff know-how into a permanent, queryable company asset.

Value creation

Owning the models lets IT automate and quality-assure process after process — the capability compounds.

Sovereignty

No enterprise data or prompts sent to external LLMs — compliance by architecture, not by contract.

Security

Cyber-assurance workloads inspect your logs, configs and tickets without any of it going offsite.

Economics

Dedicated GPUs give a flat, forecastable monthly cost instead of per-token bills that grow with every agent.

Latency & bandwidth

Inference and training beside the data keep GPU utilisation — and energy efficiency — high.

THE STACK · GPU TO USER

L5 Interfaces
Web console, assistants, agents, SQL Agent, APIs — behind your identity provider

L4 Model repository
Curated open-weight LLM / LMM plus your own, versioned and access-controlled

L3 Transformers & extensions
Serving and fine-tuning runtime, INT8 / FP8 quantisation, retrieval, connectors

L2 Clustering & orchestration
Multi-node scheduling, per-GPU telemetry, alerting, capacity planning

L1 GPU servers
HGX B300 / H200, AMD Instinct, RTX PRO — configured and tested in Australia via Hyperscalers

BENCHMARK · LOCAL VS CLOUD

Browser-automation task from natural-language instructions: configure a vehicle on a manufacturer's website and complete the enquiry form.

Qwen 2.5 72B · single NVIDIA H200

2 min 25 s

INT8, 141 GB HBM3e — whole model on one GPU

GPT-4o · public API

2 min 40 s

Data leaves the network on every call

Qwen 2.5 72B · multi-GPU A6000

6 min 22 s

Model split across cards — NCCL / PCIe overhead

Hyperscalers internal benchmark, single run, indicative only.

TRACK RECORD

Built in Australia, by the teams that now build for QORINAI.

QORINAI is a new company with an experienced bench. Its delivery partners and engineering leads have designed, built and operated high-density compute across New South Wales and Victoria since 2008. The projects are theirs; the capabilities are now ours. Counterparties and site details are available under NDA.

26 MW+
high-density compute delivered

80 MW+
liquid cooling delivered by our engineering leads

27 yrs
data-centre infrastructure experience

2
NSW distribution networks connected to

CAPABILITY TIMELINE · FROM INNER-CITY COLOCATION TO LIQUID-COOLED GPU — AND NOW GIGAWATT CAMPUSES

2008
Inner-city colocation opens

100 racks, N+1, dark-fibre ring — still operating.

2014
Open-OEM manufacturing

Our strategic shareholder begins building servers in Canberra.

2021–22
Modular site beside a power station

8 × 2.5 MW modules, 1.9 → 20 MW staged in five months.

2022
~14 MW high-density build

Largest single build in the portfolio; direct utility coordination.

2025–26
1 MW liquid-cooled GPU cluster

Owner-funded, no master EPC, PUE < 1.10 pod design.

2026
GW-class campus masterplan

QORINAI's own programme: power first, three supply layers.

Four projects, four links of the chain.

Each project covers a different capability — grid connection, modular speed, continuous operations, liquid-cooled owner-managed delivery — and the same people carry across all four.

PROJECT 1 · HIGH-DENSITY COMPUTE

~14 MW high-density facility, Greater Sydney

~14 MW	> 1 M CFM	4–12 °C
IT load	custom cooling	hall band

Energised through direct coordination with the distribution network operator in one of Australia's most connection-constrained regions. Custom high-airflow cooling, per-tenant VLAN isolation, layered physical security.

Why it matters A formal connection in a queue-constrained network is the scarcest asset in Australian DC development. This team has done it.

PROJECT 2 · MODULAR DATA CENTRE

~11 MW modular site beside a renewable power station, Northern NSW

8 × 2.5 MW	3 days	5 months
modules	signed → first power	1.9 → 20 MW staged

Eight transportable modules behind the meter of a 30 MW biomass station under a seven-year take-or-pay agreement. Later relocated overseas and still operating — the design is genuinely redeployable.

Why it matters Modular, behind-the-meter, staged energisation is exactly how QORINAI launches a campus before the transmission connection lands.

PROJECT 3 · COLOCATION OPERATIONS

Inner-city colocation facility, Sydney — operated since 2008

100	N+1	18 yrs
racks	power & cooling	continuous operation

One train stop from the CBD: N+1 power and cooling, a redundant dark-fibre ring to every major Sydney data centre and peering point, private cages, remote hands under ISO 27001 controls.

Why it matters Eighteen years of uptime is the credit basis for QORINAI's operations SLA.

PROJECT 4 · LIQUID-COOLED GPU CLUSTER

1 MW liquid-cooled GPU cluster, Sydney north shore — no master EPC

1 MW	< 1.10	0
liquid-cooled GPU	PUE, pod design	master EPC contracts

Delivered by directly engaging electrical, mechanical / liquid-cooling, cabling and compliance subcontractors. Fed a sealed high-security pod product line with patent-pending waste-heat reuse.

Why it matters HGX B300 racks exceed 100 kW and must be liquid-cooled. This team has already done it without an EPC in the middle.

Counterparties and site details available under NDA

FLAGSHIP PROJECT · MELBOURNE'S NORTHERN CORRIDOR, VICTORIA

A gigawatt AI campus designed around its power.

A single greenfield freehold beside a 500 kV transmission corridor, with a formal 1,000 MW connection response from the state's transmission planner already in hand. Built as repeatable 250 MW modules towards 1 GW of IT load, with gas firming and 500 MWh of battery storage fed from a new 500 kV terminal station — the power process was run first, so the campus can be.

1,000 MW
Maximum demand assessed — connection feasible

4 × 250 MW
Repeatable modules to 1 GW IT

500 MWh
Grid-forming battery storage on site

36 months
Application to first energisation



CAMPUS MASTERPLAN VISUALISATION · INDICATIVE

01 · 500 KV BACKBONE
The transmission line is 1 km away

02 · FIRST MOVER
One new terminal station per radius

03 · FORMAL RESPONSE
Rules-based, not a desktop enquiry

04 · SINGLE TITLE
One surveyed freehold lot, ~70 ha

05 · GAS CONFIRMED
Pipeline scoped for 300 MW of firming

06 · THE CORRIDOR
Where the hyperscalers are already going

Site identification, commercial terms and the full teaser are available under NDA

Three supply layers back each other up. 500 kV in, 220 kV across the campus, 33 kV into each building.

PRIMARY · TRANSMISSION

New 500 kV shared terminal station

Cut-in to the adjacent 500 kV line, campus supplied at 220 kV. A contestable asset built by a nominated transmission service provider; indicative commissioning window 2029.

INTERIM · DISTRIBUTION

Staged distribution pathway

4 MVA accessible today from two substations inside the boundary. The distributor has responded to a 1,000 MVA enquiry with a staged pathway and offered to lead the feasibility study.

CONTINGENT · ON-SITE

Gas firming 300 MW + 500 MWh BESS

Gas volume confirmed via a new 6–8 km pipeline. A grid-forming battery provides frequency response, black start and peak shaving — the natural remediation asset for system strength.

THE AUTHORITIES HAVE ANSWERED

Grid · transmission	1,000 MW enquiry lodged Apr 2026; formal NER 5.3.3 response Jun 2026 — 500 kV terminal station, three options feasible.	FEASIBLE · FORMAL RESPONSE
Grid · distribution	Interim supply for construction and initial IT load; distributor offered to lead the feasibility study.	STAGED PATHWAY
Gas	Volume for 300 MW of firming achievable via a new 6–8 km pipeline.	FEASIBLE · QUANTIFIED
Water · sewer · comms	Water assessment under way; sewer routes identified; all major carriers at the boundary.	FEASIBLE
Land & planning	Single surveyed freehold lot; planning pathway identified under the state's facilitation programme.	PATHWAY IDENTIFIED

PROGRAMME · 36 MONTHS TO FIRST POWER

- **NOW** — **Partner & term sheet** NDA, data room, joint technical workshop
- **M 0–2** — **Settle & lodge** Land settlement, connection application, planning lodged
- **M 2–10** — **Feasibility & design** Masterplan detail, M&E, gas / BESS / substation FEED, grid models
- **M 14** — **Construction approval** Permit, Use of System and Connection Services Agreements
- **M 14–36** — **22 months of construction** 500 kV terminal station, switchyard, BESS, Building 1
- **M 36** — **Stage 1 live** Building 1 at 250 MW inside the planner's 2029 window

Preliminary information · not an offer · details under NDA

One module, repeatable four times within the 1 GW response.

Each module is a 250 MW data-hall block with its own 220/33 kV substation — sized to sit inside the 600 MW contingency limit and the planner's staged-load profile, so every repeat follows the same engineering, approvals and supply chain.

Module 1 · 250 MW

Months 0–36. Terminal station → 500/220 kV switchyard → BESS and Building 1 substation. Interim supply and gas firming during construction.

~17,870 GPUS · 17 CLUSTERS × 128

Repeat 1 · 500 MW

After first energisation, inside the 600 MW contingency limit; 2–3 connection points.

~35,700 GPUS CUMULATIVE

Repeat 2 · 750 MW

With the third connection point and the planner's staged-load profile; 500 kV line reinforcement scope.

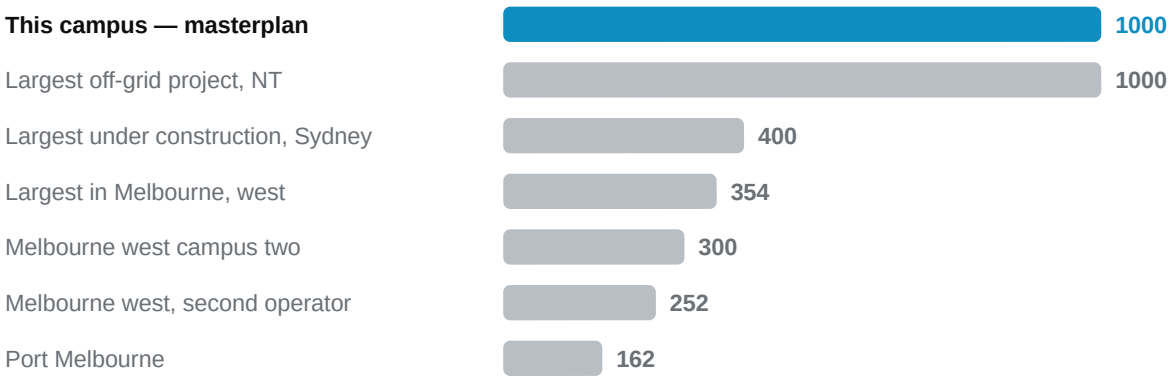
~53,600 GPUS CUMULATIVE

Repeat 3 · 1 GW

With the regional transformer upgrade (2031) and the ultimate breaker-and-a-half arrangement.

~71,500 GPUS CUMULATIVE

AUSTRALIAN BENCHMARK · CAMPUSES UNDER CONSTRUCTION, MW



Capacities as published by each operator (industry construction snapshots, 2026). Every listed grid-connected campus relies on diesel standby; this project designs gas firming and 500 MWh of BESS into the campus.

WAYS TO PARTICIPATE · COMBINABLE STRUCTURES

A · Energy partner Behind-the-meter power JV — co-own gas generation, BESS and on-site transformation under a long-term PPA or tolling agreement.

B · Equipment + equity Preferred supply of BESS, PCS, inverters and substation equipment, with minority equity in the project company.

C · Devco equity Enter before the planning application and capture the step-up from "feasible" to "approved".

D · Showcase anchor A smaller commitment for energy-system naming, an on-site demonstration centre and joint announcements.

SERVICES

Build with us, host with us, or just rent the compute.

Four service families that share one engineering team, one set of facilities and one monthly AUD invoice.

Infrastructure development

1 MW → GW

- AI data centre design & build — Tier III+ MEP, construction management, IST
- Power & grid connection — enquiries, option studies, terminal-station pathways
- Energy & microgrid — gas firming, grid-forming BESS, PPAs, take-or-pay
- Modular data centres — transportable 2.5 MW units
- Liquid cooling engineering — DLC, RDHx, CDU plant, 80–150 kW racks
- Commissioning & burn-in — load-bank IST, NCCL, HPL, thermal soak

Data centre services

30–150 KW · A+B

- High-density colocation — quarter rack to private suite
- GPU capacity hosting for owners, neoclouds and enterprise AI teams
- Private cages & suites with biometric access and separate metering
- Smart hands 24×7 — < 15 min P1, < 4 h general
- Cross-connects, IP transit, dark fibre, business fibre
- VPC & SD-WAN, DDoS scrubbing and managed firewall

Compute & cloud

AUD BILLING

- GPU compute and multi-node clusters — hourly or reserved
- GPU-as-a-Service with a fixed AUD price for the term
- AI model hosting — 100+ models, OpenAI-compatible API
- Bare metal, virtual machines, private OpenStack cloud
- Block & object storage, DR and immutable backup
- HPC and VFX rendering by the frame-hour

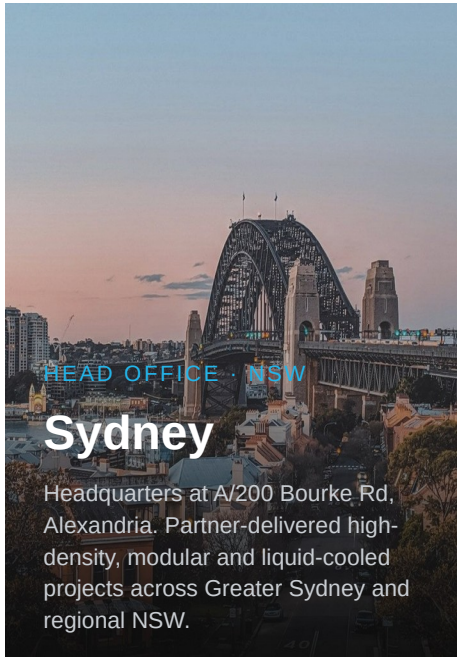
Managed & professional

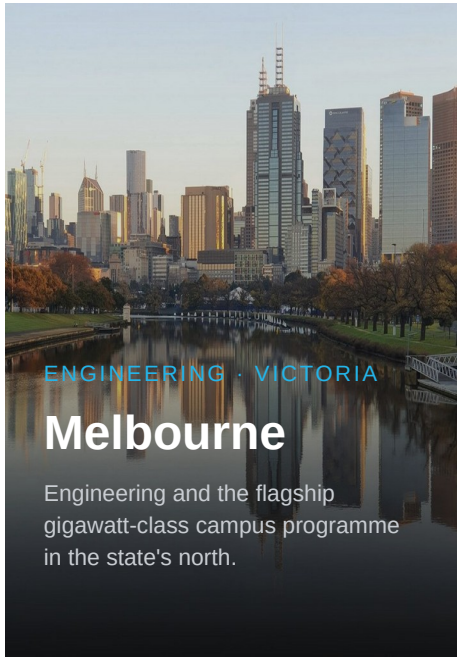
99.9 % NODE SLA

- GPU fleet management — 24×7 NOC, DCGM, firmware lifecycle, RMA
- Managed infrastructure, OS and platform layers
- Professional services — architecture, migration, cluster sizing, cost modelling
- Lab-as-a-Service — test-drive before committing capital
- IP-appliance design through the Hyperscalers channel
- 24×7 Australian support — one number, one ticket queue

WHERE WE ARE

Sydney-led, national by design.







TECHNOLOGY PLATFORMS

NVIDIA HGX B300 · GB300 NVL72 · AMD Instinct MI325X · Intel Xeon 6 · Quantum InfiniBand XDR · Spectrum-X 800G · Open Compute Project · Direct-to-chip liquid cooling · Tier III+

CREDENTIALS · STANDARDS WE OPERATE TO

Facility design	Tier III+ concurrently maintainable, 99.99 % design target
Security	ISO 27001 aligned operations and controls
Efficiency	PUE ≤ 1.2 design; partner pod tested < 1.10
Ownership	100 % Australian owned and operated

WHO WE BUILD AND RUN COMPUTE FOR

- ✓ Sovereign workloads
- ✓ AI startups & neoclouds
- ✓ GPU asset owners & lenders
- ✓ Financial services
- ✓ Enterprise
- ✓ Software & MSPs
- ✓ Research & education
- ✓ Media & VFX

QORINAI

WORK WITH QORINAI

Tell us about the megawatts and the models.

Land and a grid connection, a fleet of GPUs with nowhere to go, or a model with no compute — start at whichever link you're on. An engineer replies within one business day.

qorinai.ai · contact form, replies within one business day

A/200 Bourke Rd, Alexandria NSW 2015, Australia

START AT WHICHEVER LINK YOU'RE ON

Build an AI data centre
Design & build, power & grid, modular launch, liquid cooling

Secure power for a site
Connection enquiries, energy architecture, staged energisation

Deploy or host a GPU cluster
Hardware configured and tested in Australia, benchmarked handover, 24×7 fleet ops

Reserve sovereign compute
GPU-as-a-Service, IntraLLM, private cloud — billed in AUD

Discuss investment or partnership
Project teasers and track-record references under NDA

© 2026 QORINAI Pty Ltd · Specifications indicative and subject to final design · Track-record projects delivered by QORINAI's delivery partners; counterparties under NDA